

ALTERNATING DIRECTION METHODS FOR LATENT VARIABLE GAUSSIAN GRAPHICAL MODEL SELECTION

SHIQIAN MA ^{*}, LINGZHOU XUE [†], AND HUI ZOU [‡]

Abstract. Chandrasekaran, Parrilo and Willsky (2010) proposed a convex optimization problem to characterize graphical model selection in the presence of unobserved variables. This convex optimization problem aims to estimate an inverse covariance matrix that can be decomposed into a sparse matrix minus a low-rank matrix from sample data. Solving this convex optimization problem is very challenging, especially for large problems. In this paper, we propose two alternating direction methods for solving this problem. The first method is to apply the classical alternating direction method of multipliers to solve the problem as a consensus problem. The second method is a proximal gradient based alternating direction method of multipliers. Our methods exploit and take advantage of the special structure of the problem and thus can solve large problems very efficiently. Global convergence result is established for the proposed methods. Numerical results on both synthetic data and gene expression data show that our methods usually solve problems with one million variables in one to two minutes, and are usually five to thirty five times faster than a state-of-the-art Newton-CG proximal point algorithm.

Key words. Alternating Direction Method, Proximal Gradient, Global Convergence, Gaussian Graphical Models, Latent Variables, Sparsity, Low-rank, Regularization.

1. Introduction. In this paper, we consider alternating direction methods with the theoretical guarantee of global convergence for computing the latent-variable graphical model selection [7]. Graphical model selection is closely related to the inverse covariance matrix estimation problem, which is of fundamental importance in multivariate statistical inference. In particular, when data $X = (X_1, \dots, X_p)'$ follow a p -dimensional joint normal distribution with some unknown variance matrix Σ , the precision matrix $\Theta = \Sigma^{-1}$ can be directly translated into a Gaussian graphical model. The zero entries in the precision matrix $\Theta = (\theta_{ij})_{1 \leq i, j \leq p}$ precisely capture the desired conditional independencies in the Gaussian graphical model [31, 19], i.e. $\theta_{ij} = 0$ if and only if $X_i \perp\!\!\!\perp X_j \mid X_{-(i,j)}$. The Gaussian graphical model has been successfully used to explore complex systems consisting of Gaussian random variables in many research fields, including gene expression genomics [22, 55], image processing [33], macroeconomics determinants study [13], and social study [1, 30].

Nowadays, massive high-dimensional data are being routinely generated with rapid advances of modern high-throughput technology (e.g. microarray and functional magnetic resonance imaging). Estimation of a sparse graphical model has become increasingly important in the high-dimensional regime, and some well-developed penalization techniques have received considerable attention in the statistical literature. [40] was the first to study the high-dimensional sparse graphical model selection problem, and they proposed the neighborhood penalized regression scheme which performs the lasso [51] to fit each neighborhood regression and summarizes the sparsity pattern by aggregation via union or intersection. [44] proposed the joint sparse regression model to jointly estimate all neighborhood lasso penalized regressions. [58] considered the Dantzig selector [6] as an alternative to the lasso in each neighborhood regression. [5] proposed a constrained ℓ_1 minimization estimator called CLIME for estimating sparse precision matrices, and established rates of convergence under both the entrywise ℓ_∞ norm and the Frobenius norm. Computationally CLIME can be further decomposed into a series of vector minimization problems.

^{*}Institute for Mathematics and Its Applications, 400 Lind Hall, 207 Church Street SE, University of Minnesota, Minneapolis, MN 55455, USA. Email: maxxa007@ima.umn.edu.

[†]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08536, USA. Email: lzxue@princeton.edu

[‡]School of Statistics, University of Minnesota, Minneapolis, MN 55455, USA. Email: hzou@stat.umn.edu

The ℓ_1 -penalized maximum normal likelihood method is another popular method for graphical model selection [59, 2, 21, 47]. [47] established its rate of convergence under the Frobenius norm. Under the irrepresentable conditions, [46] obtained the convergence rates under the entrywise ℓ_∞ norm and the spectral norm. Define the entrywise ℓ_1 norm of S as the sum of absolute values of the entries of S , i.e., $\|S\|_1 := \sum_{ij} |S_{ij}|$. For a given sample covariance matrix $\hat{\Sigma} \in \mathbb{R}^{p \times p}$, the ℓ_1 -penalized maximum normal likelihood estimation can be formulated as the following convex optimization problem.

$$(1.1) \quad \min_S \langle S, \hat{\Sigma} \rangle - \log \det S + \rho \|S\|_1,$$

where the first part $\langle S, \hat{\Sigma} \rangle - \log \det S$ gives the normal log-likelihood function of S , and the entrywise ℓ_1 norm $\|S\|_1$ is used to promote the sparsity of the resulting matrix. Note that in the literature the ℓ_1 -penalized maximum normal likelihood usually uses the so-called 1-off absolute penalty $\|S\|_{1,\text{off}} := \sum_{i \neq j} |S_{ij}|$. However, $\|S\|_1$ and $\|S\|_{1,\text{off}}$ cause no difference when using our algorithm. [7] have used $\|S\|_1$ in defining their convex optimization problem and hence we follow their convention in the current paper.

The aforementioned Gaussian graphical model selection methods were proposed under the ideal setting without missing variables. The recent paper by [7] considered a more realistic scenario where the full data consist of both observed variables and missing (hidden) variables. Let $X_{p \times 1}$ be the observed variables. Suppose that there are some hidden variables $Y_{r \times 1}$ ($r \ll p$) such that (X, Y) jointly follow a multivariate normal distribution. Denote the covariance matrix by $\Sigma_{(X,Y)}$ and the precision matrix by $\Theta_{(X,Y)}$. Then we can write $\Sigma_{(X,Y)} = [\Sigma_X, \Sigma_{XY}; \Sigma_{YX}, \Sigma_Y]$ and $\Theta_{(X,Y)} = [\Theta_X, \Theta_{XY}; \Theta_{YX}, \Theta_Y]$. Given the hidden variables Y , the conditional concentration matrix of observed variables, Θ_X , is sparse for a sparse graphical model. However, the marginal concentration matrix of observed variables, $\Sigma_X^{-1} = \Theta_X - \Theta_{XY} \Theta_Y^{-1} \Theta_{YX}$, might not be a sparse matrix but a difference between the sparse term Θ_X and the low-rank term $\Theta_{XY} \Theta_Y^{-1} \Theta_{YX}$. The problem of interest is to recover the sparse conditional matrix Θ_X based on observed variables X . [7] accomplished this goal by solving a convex optimization problem under the assumption that $\Sigma_X^{-1} = S - L$ for some sparse matrix S and low-rank matrix L . The low rank assumption on L holds naturally since r is much less than p . Motivated by the success of the convex relaxation for rank-minimization problem, [7] introduced a regularized maximum normal likelihood decomposition framework called the latent variable graphical model selection (LVGLASSO) as follows.

$$(1.2) \quad \min_{S,L} \langle S - L, \hat{\Sigma}_X \rangle - \log \det(S - L) + \alpha \|S\|_1 + \beta \text{Tr}(L), \quad \text{s.t. } S - L \succ 0, L \succeq 0,$$

where $\hat{\Sigma}_X$ is the sample covariance matrix of X and $\text{Tr}(L)$ denotes the trace of matrix L . In the high-dimensional setting, [7] established the consistency theory for (1.2) concerning its recovery of the support and sign pattern of S and the rank of L .

Solving the convex optimization problem (1.2) is very challenging, especially for large problems. [7] considered (1.2) as a log-determinant semidefinite programming (SDP) problem, and used a Newton-CG based proximal point algorithm (LogdetPPA) proposed by [52] to solve it. However, LogdetPPA does not take advantage of the special structure of the problem, and we argue that it is inefficient for solving large-scale problems. To illustrate our point, let us consider the special case of (1.2) with $L = 0$, and then the latent variable graphical model selection (1.2) exactly reduces to the Gaussian graphical model selection

(1.1). Note that (1.1) can be rewritten as

$$\min_S \max_{\|W\|_\infty \leq \rho} -\log \det S + \langle \hat{\Sigma}_X + W, X \rangle,$$

where $\|W\|_\infty$ is the largest absolute value of the entries of U . The dual problem of (1.1) can be obtained by exchanging the order of max and min, i.e.,

$$\max_{\|W\|_\infty \leq \rho} \min_S -\log \det X + \langle \hat{\Sigma}_X + W, S \rangle,$$

which is equivalent to

$$(1.3) \quad \max_W \{\log \det W + p : \|W - \hat{\Sigma}_X\|_\infty \leq \rho\}.$$

Both the primal and the dual graphical Lasso problems (1.1) and (1.3) can be viewed as semidefinite programming problems and can be solved via interior point methods (IPMs) in polynomial time [4]. However, the per-iteration computational cost and memory requirements of an IPM are prohibitively high for (1.1) and (1.3), especially when the size of the matrix is large. Customized SDP based methods such as the ones studied in [52] and [32] require a reformulation of the problem that increases the size of the problem and thus makes them impractical for solving large-scale problems. Therefore, most of the methods developed for solving (1.1) and (1.3) are first-order methods. These methods include block coordinate descent type methods [2, 21, 49, 56], projected gradient method [15] and variants of Nesterov's accelerated method [11, 35]. Recently, alternating direction methods have been applied to solve (1.1) and shown to be very effective [60, 48].

In this paper, we propose two alternating direction type methods to solve the latent variable graphical model selection. The first method is to apply the alternating direction method of multipliers to solve this problem. This is due to the fact that the latent variable graphical model selection can be seen as a special case of the consensus problem discussed in [3]. The second method we propose is an alternating direction method with proximal gradient steps. To apply the second method, we first group the variables into two blocks and then apply the alternating direction method with one of the subproblems being solved inexactly by taking a proximal gradient step. Our methods exploit and take advantage of the special structure of the problem and thus can solve large problems very efficiently. Although the convergence results of the proposed methods are not very different from the existing results for alternating direction type methods, we still include the convergence proof for the second method in the appendix for completeness. We apply the proposed methods to solving problems from both synthetic data and gene expression data and show that our method outperform the state-of-the-art Newton-CG proximal point algorithm LogdetPPA significantly on both accuracy and CPU times.

The rest of this paper is organized as follows. In Section 2, we give some preliminaries on alternating direction method of multipliers and proximal mappings. In Section 3, we propose solving LVGLASSO (1.2) as a consensus problem using the classical alternating direction method of multipliers. We propose the proximal gradient based alternating direction method for solving (1.2) in Section 4. In Section 5, we apply our alternating direction method to solving (1.2) using both synthetic data and gene expression data. We draw some conclusions in Section 6.

2. Preliminaries. Problem (1.2) can be rewritten in the following equivalent form by introducing a new variable R :

$$(2.1) \quad \begin{aligned} \min \quad & \langle R, \hat{\Sigma}_X \rangle - \log \det R + \alpha \|S\|_1 + \beta \text{Tr}(L) \\ \text{s.t.} \quad & R = S - L, R \succ 0, L \succeq 0, \end{aligned}$$

which can be further reduced to

$$(2.2) \quad \min \langle R, \hat{\Sigma}_X \rangle - \log \det R + \alpha \|S\|_1 + \beta \text{Tr}(L) + \mathcal{I}(L \succeq 0), \quad \text{s.t.} \quad R - S + L = 0,$$

where the indicator function $\mathcal{I}(L \succeq 0)$ is defined as

$$(2.3) \quad \mathcal{I}(L \succeq 0) := \begin{cases} 0, & \text{if } L \succeq 0 \\ +\infty, & \text{otherwise.} \end{cases}$$

Note that we have dropped the constraint $R \succ 0$ since it is already implicitly imposed by the $\log \det R$ function.

Now since the objective function involves three separable convex functions and the constraint is simply linear, Problem (2.2) is suitable for alternating direction method of multipliers (ADMM). ADMM is closely related to the Douglas-Rachford and Peaceman-Rachford operator-splitting methods for finding zero of the sum of two monotone operators that have been studied extensively in [14, 43, 34, 16, 18, 9, 10]. ADMM has been revisited recently due to its success in the emerging applications of structured convex optimization problems arising from image processing, compressed sensing, machine learning, semidefinite programming and statistics etc. (see e.g., [24, 23, 53, 57, 27, 45, 60, 48, 25, 26, 39, 54, 3, 42, 37]).

Problem (2.2) is suitable for alternating direction methods because the three convex functions involved in the objective function, i.e.,

$$(2.4) \quad f(R) := \langle R, \hat{\Sigma}_X \rangle - \log \det R,$$

$$(2.5) \quad g(S) := \alpha \|S\|_1,$$

and

$$(2.6) \quad h(L) := \beta \text{Tr}(L) + \mathcal{I}(L \succeq 0),$$

have easy proximal mappings. Note that the proximal mapping of function $c : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ for given $\xi > 0$ and $Z \in \mathbb{R}^{m \times n}$ is defined as

$$(2.7) \quad \text{prox}(c, \xi, Z) := \underset{X \in \mathbb{R}^{m \times n}}{\text{argmin}} \frac{1}{2\xi} \|X - Z\|_F^2 + c(X).$$

The proximal mapping of $f(R)$ defined in (2.4) is

$$(2.8) \quad \text{prox}(f, \xi, Z) := \underset{R}{\text{argmin}} \frac{1}{2\xi} \|R - Z\|_F^2 + \langle R, \hat{\Sigma}_X \rangle - \log \det R.$$

The first-order optimality conditions of (2.8) are given by

$$(2.9) \quad R + \xi \hat{\Sigma}_X - Z - \xi R^{-1} = 0.$$

It is easy to verify that

$$(2.10) \quad R := U \operatorname{diag}(\gamma) U^\top,$$

satisfies (2.9) and thus gives the optimal solution of (2.8), where $U \operatorname{diag}(\sigma) U^\top$ is the eigenvalue decomposition of matrix $\xi \hat{\Sigma}_X - Z$ and

$$(2.11) \quad \gamma_i = \left(-\sigma_i + \sqrt{\sigma_i^2 + 4\xi} \right) / 2, \forall i = 1, \dots, p.$$

Note that (2.11) guarantees that the solution of (2.8) given by (2.10) is a positive definite matrix. The proximal mapping of $g(S)$ defined in (2.5) is

$$(2.12) \quad \operatorname{prox}(g, \xi, Z) := \operatorname{argmin}_S \frac{1}{2\xi} \|S - Z\|_F^2 + \alpha \|S\|_1.$$

It is well known that (2.12) has a closed-form solution that is given by the ℓ_1 shrinkage operation

$$S^{k+1} := \operatorname{Shrink}(Z, \alpha\xi),$$

where $\operatorname{Shrink}(\cdot, \cdot)$ is defined as

$$(2.13) \quad [\operatorname{Shrink}(Z, \tau)]_{ij} := \begin{cases} Z_{ij} - \tau, & \text{if } Z_{ij} > \tau \\ Z_{ij} + \tau, & \text{if } Z_{ij} < -\tau \\ 0, & \text{if } -\tau \leq Z_{ij} \leq \tau. \end{cases}$$

The proximal mapping of $h(L)$ defined in (2.6) is

$$(2.14) \quad \operatorname{prox}(h, \xi, Z) := \operatorname{argmin}_L \frac{1}{2\xi} \|L - Z\|_F^2 + \beta \mathbf{Tr}(L) + \mathcal{I}(L \succeq 0).$$

It is easy to verify that the solution of (2.14) is given by

$$(2.15) \quad L := U \operatorname{diag}(\gamma) U^\top,$$

where $Z = U \operatorname{diag}(\sigma) U^\top$ is the eigenvalue decomposition of Z and γ is given by

$$(2.16) \quad \gamma_i := \max\{\sigma_i - \xi\beta, 0\}, \quad i = 1, \dots, p.$$

Note that (2.16) guarantees that L given in (2.15) is a positive semidefinite matrix.

The discussions above suggest the following natural ADMM for solving (2.2) be efficient.

$$(2.17) \quad \begin{cases} R^{k+1} &:= \operatorname{argmin}_R \mathcal{L}_\mu(R, S^k, L^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \mathcal{L}_\mu(R^{k+1}, S, L^k; \Lambda^k) \\ L^{k+1} &:= \operatorname{argmin}_L \mathcal{L}_\mu(R^{k+1}, S^{k+1}, L; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - (R^{k+1} - S^{k+1} + L^{k+1})/\mu, \end{cases}$$

where the augmented Lagrangian function is defined as

$$(2.18) \quad \mathcal{L}_\mu(R, S, L; \Lambda) := \langle R, \hat{\Sigma}_X \rangle - \log \det R + \alpha \|S\|_1 + \beta \operatorname{Tr}(L) + \mathcal{I}(L \succeq 0) - \langle \Lambda, R - S + L \rangle + \frac{1}{2\mu} \|R - S + L\|_F^2,$$

Λ is the Lagrange multiplier and $\mu > 0$ is the penalty parameter. Note that the three subproblems in (2.17) correspond to the proximal mappings of f , g and h defined in (2.4), (2.5) and (2.6), respectively. Thus they are all easy to solve. However, the global convergence of ADMM (2.17) with three blocks of variables was ambiguous. Only until very recently, was it shown that (2.17) globally converges under certain conditions (see [36]). It should be noted, however, that the error bound condition required in [36] is strong and only a few classes of convex function are known that satisfy this condition.

3. ADMM for Solving (2.2) as a Consensus Problem. Problem (2.2) can be rewritten as a convex minimization problem with two blocks of variables and two separable functions as follows:

$$(3.1) \quad \begin{aligned} \min \quad & \phi(X) + \psi(Z), \\ \text{s.t.} \quad & X - Z = 0, \end{aligned}$$

where $X = (R, S, L)$, $Z = (\tilde{R}, \tilde{S}, \tilde{L})$, and

$$\phi(X) := f(R) + g(S) + h(L), \quad \psi(Z) = \mathcal{I}(\tilde{R} - \tilde{S} + \tilde{L} = 0),$$

with f, g and h defined in (2.4), (2.5) and (2.6), respectively. The ADMM applied to solving (3.1) can be described as follows:

$$(3.2) \quad \begin{cases} X^{k+1} &:= \operatorname{argmin}_X \phi(X) - \langle \Lambda^k, X - Z^k \rangle + \frac{1}{2\mu} \|X - Z^k\|_F^2, \\ Z^{k+1} &:= \operatorname{argmin}_Z \psi(Z) - \langle \Lambda^k, X^{k+1} - Z \rangle + \frac{1}{2\mu} \|X^{k+1} - Z\|_F^2, \\ \Lambda^{k+1} &:= \Lambda^k - (X^{k+1} - Z^{k+1})/\mu, \end{cases}$$

where Λ is the Lagrange multiplier associated with the equality constraint. The two subproblems in (3.2) are both easy to solve. In fact, the solution of the first subproblem in (3.2) corresponds to the proximal mappings of f, g and h . Partitioning the matrix $T^k := X^{k+1} - \mu\Lambda^k$ into three blocks in the same form as $Z = (\tilde{R}, \tilde{S}, \tilde{L})$, The second subproblem can be reduced to:

$$(3.3) \quad \begin{aligned} \min \quad & \frac{1}{2} \|(\tilde{R}, \tilde{S}, \tilde{L}) - (T_R^k, T_S^k, T_L^k)\|_F^2 \\ \text{s.t.} \quad & \tilde{R} - \tilde{S} + \tilde{L} = 0. \end{aligned}$$

The first-order optimality conditions of (3.3) are given by

$$(3.4) \quad (\tilde{R}, \tilde{S}, \tilde{L}) - (T_R^k, T_S^k, T_L^k) - (\Gamma, -\Gamma, \Gamma) = 0,$$

where Γ is the Lagrange multiplier associated with (3.3). Thus we get,

$$\tilde{R} = T_R^k + \Gamma, \quad \tilde{S} = T_S^k - \Gamma, \quad \tilde{L} = T_L^k + \Gamma.$$

Substituting them into the equality constraint in (3.3), we get

$$(3.5) \quad \Gamma = -(T_R^k - T_S^k + T_L^k)/3.$$

By substituting (3.5) into (3.4) we get the solution to (3.3).

The ADMM (3.2) solves Problem (3.1) with two blocks of variables. It can be seen as a special case of the consensus problem discussed in [3]. The global convergence result of (3.2) has also been well studied in the literature (see e.g., [16, 18]).

4. A Proximal Gradient based Alternating Direction Method. In this section, we propose another alternating direction type method to solve (2.2). In Section 3, we managed to reduce the original problem with three blocks of variables (2.2) to a new problem with two blocks of variables (3.1). As a result, we can use ADMM for solving problems with two blocks of variables, whose convergence has been well studied. Another way to reduce the problem (2.2) into a problem with two blocks of variables is to group two variables (say S and L) as one variable. This leads to the new equivalent form of (2.2):

$$(4.1) \quad \begin{aligned} \min \quad & f(R) + \varphi(W) \\ \text{s.t.} \quad & R - [I, -I]W = 0, \end{aligned}$$

where $W = [S; L]$ and $\varphi(W) = g(S) + h(L)$. Now the ADMM for solving (3.1) can be described as

$$(4.2) \quad \begin{cases} R^{k+1} &:= \operatorname{argmin}_R f(R) - \langle \Lambda^k, R - [I, -I]W^k \rangle + \frac{1}{2\mu} \|R - [I, -I]W^k\|_F^2 \\ W^{k+1} &:= \operatorname{argmin}_W \varphi(W) - \langle \Lambda^k, R^{k+1} - [I, -I]W \rangle + \frac{1}{2\mu} \|R^{k+1} - [I, -I]W\|_F^2 \\ \Lambda^{k+1} &:= \Lambda^k - (R^{k+1} - [I, -I]W^{k+1})/\mu, \end{cases}$$

where Λ is the Lagrange multiplier associated with the equality constraint and $\mu > 0$ is a penalty parameter. The first subproblem in (4.2) is still easy and it corresponds to the proximal mapping of function f . However, the second subproblem in (4.2) is not easy, because the two parts of W are coupled together in the quadratic penalty term. To overcome this difficulty, we solve the second subproblem in (4.2) inexactly by one step of a proximal gradient method. Note that the second subproblem in (4.2) can be reduced to

$$(4.3) \quad W^{k+1} := \operatorname{argmin}_W \varphi(W) + \frac{1}{2\mu} \|R^{k+1} - [I, -I]W - \mu\Lambda^k\|_F^2.$$

One step of proximal gradient method solves the following problem

$$(4.4) \quad \min_W \varphi(W) + \frac{1}{2\mu\tau} \|W - (W^k + \tau \begin{pmatrix} I \\ -I \end{pmatrix} (R^{k+1} - [I, -I]W^k - \mu\Lambda^k))\|_F^2.$$

Since the two parts of $W = [S; L]$ are separable in the quadratic part now, (4.4) reduces to two problems

$$(4.5) \quad \min_S g(S) + \frac{1}{2\mu\tau} \|S - (S^k + \tau G_R^k)\|_F^2,$$

and

$$(4.6) \quad \min_L h(L) + \frac{1}{2\mu\tau} \|L - (L^k - \tau G_R^k)\|_F^2,$$

where $G_R^k = R^{k+1} - S^k + L^k - \mu\Lambda^k$. Both (4.5) and (4.6) are easy to solve as they correspond to the proximal mappings of functions g and h , respectively. Thus, our proximal gradient based alternating direction method (PGADM) can be summarized as

Algorithm 1 A Proximal Gradient based Alternating Direction Method

```

1: for  $k=0,1,\dots$  do
2:    $R^{k+1} := \operatorname{argmin}_R f(R) - \langle \Lambda^k, R - S^k + L^k \rangle + \frac{1}{2\mu} \|R - S^k + L^k\|_F^2$ 
3:    $S^{k+1} := \operatorname{argmin}_S g(S) + \frac{1}{2\mu\tau} \|S - (S^k + \tau G_R^k)\|_F^2$ 
4:    $L^{k+1} := \operatorname{argmin}_L h(L) + \frac{1}{2\mu\tau} \|L - (L^k - \tau G_R^k)\|_F^2$ 
5:    $\Lambda^{k+1} := \Lambda^k - (R^{k+1} - S^{k+1} + L^{k+1})/\mu$ 
6: end for

```

REMARK 4.1. *The idea of incorporating proximal step into the alternating direction method of multipliers has been suggested by [17] and [8]. This idea has then been generalized by [28] to allow varying penalty and proximal parameters. Recently, this technique has been used for sparse and low-rank optimization problems (see [57] and [50]). More recently, some convergence properties of alternating direction methods with proximal gradient steps have been studied by [12], [36], [20] and [38]. However, for the seek of completeness, we include a global convergence proof for Algorithm 1 in the Appendix.*

REMARK 4.2. *In Algorithm 1, we grouped S and L as one block of variable. We also implemented the other two ways of grouping the variables, i.e., group R and S as one block, and group R and L as one block. We found from the numerical experiments that these two alternatives yielded similar practical performance as Algorithm 1.*

REMARK 4.3. *If we use the 1-off absolute penalty $\|S\|_{1,\text{off}} := \sum_{i \neq j} |S_{ij}|$ to replace $\|S\|_1$, our algorithm basically remains the same except that we modify $\text{Shrink}(\cdot, \cdot)$ as follows*

$$(4.7) \quad [\text{Shrink}(Z, \tau)]_{ij} := \begin{cases} Z_{ii}, & \text{if } i = j \\ Z_{ij} - \tau, & \text{if } i \neq j \text{ and } Z_{ij} > \tau \\ Z_{ij} + \tau, & \text{if } i \neq j \text{ and } Z_{ij} < -\tau \\ 0, & \text{if } i \neq j \text{ and } -\tau \leq Z_{ij} \leq \tau. \end{cases}$$

5. Numerical experiments. In this section, we present numerical results on both synthetic and real data to demonstrate the efficiency of the proposed methods: ADMM (3.2) and PGADM (Algorithm 1). Our codes were written in MATLAB. All numerical experiments were run in MATLAB 7.12.0 on a laptop with Intel Core I5 2.5 GHz CPU and 4GB of RAM.

We first compared ADMM (3.2) with PGADM (Algorithm 1) on some synthetic problems. We compared ADMM and PGADM using two different ways of choosing μ . One set of comparisons used a fixed $\mu = 10$, and the other set of comparisons used a continuation scheme to dynamically change μ . The continuation scheme we used was to set the initial value of μ as the size of the matrix p , and then multiply μ by 1/4 after every 10 iterations.

We then compared the performance of PGADM (with continuation on μ) with LogdetPPA proposed by [52] and used in [7] for solving (2.2).

5.1. Comparison of ADMM and PGADM on Synthetic Data. We observed from the numerical experiments that the step size τ of the proximal gradient step in PGADM (Algorithm 1) can be slightly larger than 1/2 and the algorithm produced very good results. We thus chose the step size τ to be 0.6 in our experiments.

We randomly created test problems using a procedure proposed by [49] and [48] for the classical graphical lasso problems. Similar procedures were used by [52] and [32]. For a given number of observed variables p and a given number of latent variables p_h , we first created a sparse matrix $U \in \mathbb{R}^{(p+p_h) \times (p+p_h)}$ with sparsity around 10%, i.e., 10% of the entries are nonzeros. The nonzero entries were set to -1 or 1 with equal probability. Then we computed $K := (U * U^\top)^{-1}$ as the true covariance matrix. We then chose the submatrix of K , $\hat{S} := K(1:p, 1:p)$ as the ground truth matrix of the sparse matrix S and chose $\hat{L} := K(1:p, p+1:p+p_h)K(p+1:p+p_h, p+1:p+p_h)^{-1}K(p+1:p+p_h, 1:p)$ as the ground truth matrix of the low rank matrix L . We then drew $N = 5p$ iid vectors, Y_1, \dots, Y_N , from the Gaussian distribution $\mathcal{N}(\mathbf{0}, (\hat{S} - \hat{L})^{-1})$ by using the *mvnrnd* function in MATLAB, and computed a sample covariance matrix of the observed variables $\Sigma_X := \frac{1}{N} \sum_{i=1}^N Y_i Y_i^\top$.

We computed the relative infeasibility of the sequence (R^k, S^k, L^k) generated by inexact ADMM using

$$(5.1) \quad \text{infeas} := \frac{\|R^k - S^k + L^k\|_F}{\max\{1, \|R^k\|_F, \|S^k\|_F, \|L^k\|_F\}}.$$

In the comparison of ADMM and PGADM, the size of all problems was chosen as $p = 1000$. For fixed $\mu = 10$, we first ran the ADMM for 100 iterations, and recorded the objective function value and *infeas*. We then ran PGADM until it achieves an objective function value within relative error 10^{-5} compared with the objective function value given by ADMM, or it achieves an *infeas* within relative error 10^{-5} compared with the *infeas* given by ADMM. The number of iterations, CPU times, *infeas* and objective function values for both ADMM and PGADM were reported in Table 5.1. From Table 5.1, we see that for fixed $\mu = 10$, ADMM was faster than PGADM when α and β are both small, and PGADM was faster than ADMM when α and β are both large.

TABLE 5.1
Comparison of ADMM with PGADM (for fixed μ) on synthetic data

		PGADM				ADMM			
α	β	obj	iter	cpu	infeas	obj	iter	cpu	infeas
0.005	0.025	-1.6987e+002	135	224.1	7.6e-005	-1.7098e+002	100	172.9	6.6e-005
0.005	0.05	-9.3385e+001	146	243.7	3.8e-005	-9.3476e+001	100	173.7	2.9e-005
0.01	0.05	-4.4748e+001	132	214.7	1.7e-005	-4.4761e+001	100	180.1	7.0e-006
0.01	0.1	5.4571e+001	111	177.4	9.8e-006	5.4567e+001	100	166.9	2.8e-007
0.02	0.1	1.1881e+002	83	136.8	9.7e-006	1.1881e+002	100	170.1	1.9e-007
0.02	0.2	2.3717e+002	56	91.5	9.3e-006	2.3717e+002	100	174.0	1.6e-007
0.04	0.2	3.2417e+002	44	67.0	9.2e-006	3.2417e+002	100	165.2	1.6e-007
0.04	0.4	4.5701e+002	19	29.4	3.1e-004	4.5700e+002	100	168.4	1.7e-007

We then further compare ADMM and PGADM on synthetic data with the continuation scheme for μ discussed above. We terminated both ADMM and PGADM when *infeas* $< 10^{-5}$. We reported the results in Table 5.2.

From Table 5.2, we see that the continuation scheme used really helped to speed up the convergence and produced much better results. Also, using this continuation scheme, PGADM was faster than ADMM with

comparable residuals and objective function values. However, we should remark that PGADM was faster than ADMM using the specific continuation scheme. If other continuation schemes were adopted, the results could be quite different. In the comparison with LogdetPPA in the following sections, we only compare LogdetPPA with PGADM with this continuation scheme.

TABLE 5.2
Comparison of ADMM with PGADM (with continuation for μ) on synthetic data

		PGADM				ADMM			
α	β	obj	iter	cpu	resid	obj	iter	cpu	resid
0.005	0.025	-1.711329e+002	32	51.5	5.7e-006	-1.711334e+002	61	103.7	6.7e-006
0.005	0.05	-9.348245e+001	41	65.8	4.2e-006	-9.348589e+001	62	103.2	6.6e-006
0.010	0.05	-4.476323e+001	41	67.8	2.8e-006	-4.476499e+001	62	105.7	9.8e-006
0.010	0.10	5.456790e+001	41	65.7	5.1e-006	5.456724e+001	71	119.6	6.0e-006
0.020	0.10	1.188077e+002	41	64.9	3.4e-006	1.188054e+002	71	115.1	8.1e-006
0.020	0.20	2.371659e+002	45	76.4	8.7e-006	2.371687e+002	75	125.6	7.0e-006
0.040	0.20	3.241688e+002	44	72.7	8.3e-006	3.241684e+002	75	126.6	8.5e-006
0.040	0.40	4.570019e+002	50	77.5	7.5e-006	4.570058e+002	78	126.5	9.7e-006

5.2. Comparison of PGADM and LogdetPPA on Synthetic Data. In this section, we compare PGADM with LogdetPPA on synthetic data created the same way as in the last section. LogdetPPA, proposed by Wang et al. in [52], is a proximal point algorithm for solving semidefinite programming problems with $\log \det(\cdot)$ function. The specialized MATLAB codes of LogdetPPA for solving (2.2) were downloaded from <http://ssg.mit.edu/~venkatc/latent-variable-code.html>.

We compared PGADM (with continuation on μ) with LogdetPPA with different α and β . We reported the comparison results on objective function value, CPU time, sparsity of S and infeas in Table 5.3. The sparsity of S is denoted as

$$sp := \frac{\#\{(i, j) : S_{ij} \neq 0\}}{p^2},$$

i.e., the percentage of nonzero entries. Since matrix S generated by LogdetPPA is always dense but with many small entries, we also measure its sparsity by truncating small entries that less than 10^{-4} to zeros, i.e.,

$$sp1 := \frac{\#\{(i, j) : |S_{ij}| > 10^{-4}\}}{p^2}.$$

All CPU times reported are in seconds. We report the speed up of PGADM over LogdetPPA in Table 5.4.

From Table 5.3 we see that the solutions produced by PGADM always have comparable objective function values compared to the solutions produced by LogdetPPA. However, our PGADM is always much faster than LogdetPPA, as shown in both Tables 5.3 and 5.4. In fact, PGADM is usually ten times faster than LogdetPPA, and sometimes more than thirty five times faster. For example, for the four large problems with matrices size 2000×2000 , LogdetPPA needs 1 hour 23 minutes, 1 hour 26 minutes, 1 hour 47 minutes and 1 hour 20 minutes, respectively, to solve them, while our PGADM needs about 7 minutes, 7 minutes, 8 minutes and 9 minutes respectively to solve them. We also notice that the matrix S generated by PGADM is always a sparse matrix with many entries exactly equal to zero, but S generated by LogdetPPA is always a dense matrix, and only when we truncate the entries that are smaller than 10^{-4} to zeros, it becomes a sparse matrix with similar level of sparsity. This is because in our PGADM, S is updated by the ℓ_1 shrinkage

TABLE 5.3
Results of PGADM and LogdetPPA on synthetic data

dim	LogdetPPA				PGADM			
p	obj	cpu	sp (%)	sp1 (%)	obj	cpu	sp (%)	infeas
$\alpha = 0.005, \beta = 0.025$								
200	1.914315e+2	7.2	100.00	19.17	1.910379e+2	1.0	18.90	4.3e-6
500	1.898418e+2	235.2	100.00	5.78	1.898275e+2	7.3	5.63	7.2e-6
1000	-1.711293e+2	1706.0	100.00	0.52	-1.711329e+2	48.2	0.49	5.7e-6
2000	-1.430010e+3	5001.2	100.00	0.06	-1.435605e+3	427.2	0.05	4.9e-6
$\alpha = 0.005, \beta = 0.05$								
200	1.926376e+2	29.1	100.00	43.63	1.924829e+2	2.5	48.66	6.6e-6
500	2.051884e+2	358.3	100.00	12.04	2.051425e+2	10.5	11.42	8.8e-6
1000	-9.347297e+1	1076.4	100.00	4.88	-9.348245e+1	71.6	4.72	4.2e-6
2000	-1.229323e+3	5191.5	100.00	0.21	-1.230238e+3	445.0	0.15	9.5e-6
$\alpha = 0.01, \beta = 0.05$								
200	2.030390e+2	20.8	100.00	20.46	2.026586e+2	2.6	11.06	9.0e-6
500	2.394720e+2	146.0	100.00	4.25	2.394631e+2	10.7	4.14	3.7e-6
1000	-4.476078e+1	740.6	100.00	0.24	-4.476323e+1	72.8	0.23	2.8e-6
2000	-1.101454e+3	6433.4	100.00	0.05	-1.111504e+3	453.7	0.05	6.4e-6
$\alpha = 0.01, \beta = 0.1$								
200	2.050879e+2	29.1	100.00	42.16	2.048359e+2	1.6	35.83	7.0e-6
500	2.639548e+2	235.2	100.00	8.62	2.638565e+2	11.6	8.19	5.8e-6
1000	5.456825e+1	932.1	100.00	2.26	5.456790e+1	74.9	2.26	5.1e-6
2000	-8.541802e+2	4813.6	100.00	0.14	-8.712916e+2	516.6	0.07	8.6e-6

operation, which truncates the small entries to zeros, while LogdetPPA needs to replace $\|S\|_1$ with smooth linear function which does not preserve sparsity.

TABLE 5.4
Speed up of PGADM over LogdetPPA on synthetic data.

p	$\alpha = 0.005, \beta = 0.025$	$\alpha = 0.005, \beta = 0.05$	$\alpha = 0.01, \beta = 0.05$	$\alpha = 0.01, \beta = 0.1$
200	7.1	11.7	7.9	18.7
500	32.0	34.2	13.6	20.3
1000	35.4	15.0	10.2	12.4
2000	11.7	11.7	14.2	9.3

5.3. Comparison of PGADM and LogdetPPA on Gene expression data. To further demonstrate the efficacy of PGADM, we applied PGADM to solving (2.2) with two gene expression data sets. One data set is the Rosetta Inpharmatics Compendium of gene expression data (denoted as Rosetta) [29] profiles which contains 301 samples with 6316 variables (genes). The other data set is the Iconix microarray data set (denoted as Iconix) from drug treated rat livers [41] which contains 255 samples with 10455 variables.

For a given number of observed variables p , we created the sample covariance matrix Σ_X by the following procedure. We first computed the variances of all of variables using all the sample data. We then selected the p variables with the highest variances and computed the sample covariance matrix Σ_X of these p variables using all the sample data. We reported the comparison results of PGADM and LogdetPPA in Tables 5.5 and 5.6 for the Rosetta and Iconix data sets, respectively. Table 5.7 summarizes the speed up of PGADM over LogdetPPA.

TABLE 5.5
Results of PGADM and LogdetPPA on Rosetta data set

dim	LogdetPPA				PGADM			
p	obj	cpu	sp (%)	sp1 (%)	obj	cpu	sp (%)	infeas
200	-2.726188e+2	5.6	100.00	0.58	-2.726184e+2	0.7	0.58	7.1e-6
500	-8.662116e+2	68.0	100.00	0.20	-8.662113e+2	7.9	0.20	3.7e-6
1000	-1.970974e+3	490.9	100.00	0.10	-1.970973e+3	52.2	0.10	7.7e-6
2000	-4.288406e+3	4597.9	100.00	0.05	-4.288406e+3	422.3	0.05	5.4e-6

TABLE 5.6
Results of PGADM and LogdetPPA on Iconix data set

dim	LogdetPPA				PGADM			
p	obj	cpu	sp (%)	sp1 (%)	obj	cpu	sp (%)	infeas
200	1.232884e+3	38.5	100.00	36.10	1.232744e+3	2.5	41.62	7.2e-6
500	1.842623e+3	98.8	100.00	2.27	1.839838e+3	17.0	0.77	1.0e-5
1000	1.439052e+3	1341.9	100.00	1.45	1.435425e+3	94.6	0.14	6.3e-6
2000	1.242966e+2	13207.2	100.00	0.07	1.168757e+2	738.2	0.06	8.5e-6

From Table 5.5, we again see that PGADM always generates solutions with comparable objective function values in much less time. For example, for $p = 2000$, LogdetPPA needs 1 hour 16 minutes to solve it while PGADM takes just 7 minutes. From Table 5.6, we see that the advantage of PGADM is more obvious. For $p = 200, 500, 1000$ and 2000 , PGADM always generates solutions with much smaller objective function values and it is always much faster than LogdetPPA. For example, for $p = 2000$, LogdetPPA takes 3 hours 40 minutes to solve it while PGADM just takes about 12 minutes.

TABLE 5.7
Speed up of PGADM over LogdetPPA on Rosetta and Iconix data sets.

dim	Rosetta data	Iconix data
200	8.0	15.4
500	8.6	5.8
1000	9.4	14.2
2000	10.9	17.9

6. Conclusion. In this paper, we proposed alternating direction methods for solving latent variable Gaussian graphical model selection. The global convergence results of our methods were established. We applied the proposed methods for solving large problems from both synthetic data and gene expression data. The numerical results indicated that our methods were five to thirty five times faster than a state-of-the-art Newton-CG proximal point algorithm.

Acknowledgement. The authors thank Professor Stephen Boyd for suggesting solving (2.2) as a consensus problem (3.1) using ADMM. Shiqian Ma's research is supported by the National Science Foundation postdoctoral fellowship through the Institute for Mathematics and Its Applications at University of Minnesota. Hui Zou's research is supported in part by grants from the National Science Foundation and the Office of Naval Research.

- [1] A. AHMED AND E.P. XING, *Recovering time-varying networks of dependencies in social and biological studies*, Proceedings of the National Academy of Sciences, 106 (2009), pp. 11878–11883.
- [2] O. BANERJEE, L. EL GHOU, AND A. D’ASPREMONT, *Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data*, Journal of Machine Learning Research, 9 (2008), pp. 485–516.
- [3] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Foundations and Trends in Machine Learning, (2011).
- [4] S. BOYD AND L. VANDENBERGHE, *Convex optimization*, Cambridge University Press, Cambridge, 2004.
- [5] T. CAI, W. LIU, AND X. LUO, *A Constrained ℓ_1 Minimization Approach to Sparse Precision Matrix Estimation*, Journal of the American Statistical Association, 106 (2011), pp. 594–607.
- [6] E. CANDÈS AND T. TAO, *The Dantzig selector: Statistical estimation when p is much larger than n* , The Annals of Statistics, 35 (2007), pp. 2313–2351.
- [7] V. CHANDRASEKARAN, P.A. PARRILO, AND A.S. WILLSKY, *Latent variable graphical model selection via convex optimization*, preprint, (2010).
- [8] G. CHEN AND M. TEBoulLE, *A proximal-based decomposition method for convex minimization problems*, Mathematical Programming, 64 (1994), pp. 81–101.
- [9] P. L. COMBETTES AND JEAN-CHRISTOPHE PESQUET, *A Douglas-Rachford splitting approach to nonsmooth convex variational signal recovery*, IEEE Journal of Selected Topics in Signal Processing, 1 (2007), pp. 564–574.
- [10] P. L. COMBETTES AND V. R. WAJS, *Signal recovery by proximal forward-backward splitting*, SIAM Journal on Multiscale Modeling and Simulation, 4 (2005), pp. 1168–1200.
- [11] A. D’ASPREMONT, O. BANERJEE, AND L. EL GHOU, *First-order methods for sparse covariance selection*, SIAM Journal on Matrix Analysis and its Applications, 30 (2008), pp. 56–66.
- [12] W. DENG AND W. YIN, *On the global and linear convergence of the generalized alternating direction method of multipliers*, tech. report, Rice University CAAM, Technical Report TR12-14, 2012.
- [13] A. DOBRA, T.S. EICHER, AND A. LENKOSKI, *Modeling uncertainty in macroeconomic growth determinants using Gaussian graphical models*, Statistical Methodology, 7 (2009), pp. 292–306.
- [14] J. DOUGLAS AND H. H. RACHFORD, *On the numerical solution of the heat conduction problem in 2 and 3 space variables*, Transactions of the American Mathematical Society, 82 (1956), pp. 421–439.
- [15] J. DUCHI, S. GOULD, AND D. KOLLER, *Projected subgradient methods for learning sparse Gaussian*, Conference on Uncertainty in Artificial Intelligence (UAI 2008), (2008).
- [16] J. ECKSTEIN, *Splitting methods for monotone operators with applications to parallel optimization*, PhD thesis, Massachusetts Institute of Technology, 1989.
- [17] ———, *Some saddle-function splitting methods for convex programming*, Optimization Methods and Software, 4 (1994), pp. 75–83.
- [18] J. ECKSTEIN AND D. P. BERTSEKAS, *On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators*, Math. Program., 55 (1992), pp. 293–318.
- [19] D. EDWARDS, *Introduction to graphical modelling*, Springer Verlag, 2000.
- [20] M. FAZEL, T. PONG, D. SUN, AND P. TSENG, *Hankel matrix rank minimization with applications to system identification and realization*, preprint, (2012).
- [21] J. FRIEDMAN, T. HASTIE, AND R. TIBSHIRANI, *Sparse inverse covariance estimation with the graphical lasso*, Biostatistics, 9 (2008), pp. 432–441.
- [22] N. FRIEDMAN, *Inferring cellular networks using probabilistic graphical models*, Science, 303 (2004), p. 799.
- [23] D. GABAY, *Applications of the method of multipliers to variational inequalities*, in Augmented Lagrangian Methods: Applications to the Solution of Boundary Value Problems, M. Fortin and R. Glowinski, eds., North-Holland, Amsterdam, 1983.
- [24] R. GLOWINSKI AND P. LE TALLEC, *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*, SIAM, Philadelphia, Pennsylvania, 1989.
- [25] D. GOLDFARB AND S. MA, *Fast multiple splitting algorithms for convex optimization*, SIAM Journal on Optimization, 22 (2012), pp. 533–556.
- [26] D. GOLDFARB, S. MA, AND K. SCHEINBERG, *Fast alternating linearization methods for minimizing the sum of two convex functions*, Mathematical Programming Series A, to appear, (2012).
- [27] T. GOLDSTEIN AND S. OSHER, *The split Bregman method for L_1 -regularized problems*, SIAM J. Imaging Sci., 2 (2009), pp. 323–343.
- [28] B. S. HE, L.-Z. LIAO, D. HAN, AND H. YANG, *A new inexact alternating direction method for monotone variational inequalities*, Math. Program., 92 (2002), pp. 103–118.

- [29] T. R. HUGHES, M. J. MARTON, A. R. JONES, C. J. ROBERTS, R. STOUGHTON, C. D. ARMOUR, H. A. BENNETT, E. COFFEY, H. DAI, Y. D. HE, M. J. KIDD, A. M. KING, M. R. MEYER, D. SLADE, P. Y. LUM, S. B. STEPANIANTS, D. D. SHOEMAKER, D. GACHOTTE, K. CHAKRABURTTY, J. SIMON, M. BARD, AND S. H. FRIEND, *Functional discovery via a compendium of expression profiles*, Cell, 102 (2000), p. 109126.
- [30] M. KOLAR, L. SONG, A. AHMED, AND E.P. XING, *Estimating time-varying networks*, The Annals of Applied Statistics, 4 (2010), pp. 94–123.
- [31] S.L. LAURITZEN, *Graphical models*, Oxford University Press, USA, 1996.
- [32] L. LI AND K.-C. TOH, *An inexact interior point method for l_1 -regularized sparse covariance selection*, preprint, (2010).
- [33] S.Z. LI, *Markov random field modeling in image analysis*, Springer-Verlag New York Inc, 2009.
- [34] P. L. LIONS AND B. MERCIER, *Splitting algorithms for the sum of two nonlinear operators*, SIAM Journal on Numerical Analysis, 16 (1979), pp. 964–979.
- [35] Z. LU, *Smooth optimization approach for sparse covariance selection*, SIAM J. Optim., 19 (2009), pp. 1807–1827.
- [36] Z. Q. LUO, *On the linear convergence of alternating direction method of multipliers*, preprint, (2012).
- [37] S. MA, *Alternating direction method of multipliers for sparse principal component analysis*, preprint, (2011).
- [38] ———, *Alternating proximal gradient method for convex minimization*, preprint, (2012).
- [39] J. MALICK, J. POVH, F. RENDL, AND A. WIEGELE, *Regularization methods for semidefinite programming*, SIAM Journal on Optimization, 20 (2009), pp. 336–356.
- [40] N. MEINSHAUSEN AND P. BÜHLMANN, *High-dimensional graphs and variable selection with the lasso*, The Annals of Statistics, (2006), pp. 1436–1462.
- [41] G. NATSOULIS, L. EL GHAOU, G. LANCKRIET, A. TOLLEY, F. LEROY, S. DUNLEA, B. EYNON, C. PEARSON, S. TUGENDREICH, AND K. JARNAGIN, *Classification of a large microarray data set: algorithm comparison and analysis of drug signatures*, Genome Research, 15 (2005), p. 724 736.
- [42] N. PARIKH AND S. BOYD, *Block splitting for large-scale distributed learning*, in NIPS, 2011.
- [43] D. H. PEACEMAN AND H. H. RACHFORD, *The numerical solution of parabolic elliptic differential equations*, SIAM Journal on Applied Mathematics, 3 (1955), pp. 28–41.
- [44] J. PENG, P. WANG, N. ZHOU, AND J. ZHU, *Partial correlation estimation by joint sparse regression models*, Journal of the American Statistical Association, 104 (2009), pp. 735–746.
- [45] Z. QIN, D. GOLDFARB, AND S. MA, *An alternating direction method for total variation denoising*, preprint, (2011).
- [46] P. RAVIKUMAR, M.J. WAINWRIGHT, G. RASKUTTI, AND B. YU, *High-dimensional covariance estimation by minimizing l_1 -penalized log-determinant divergence*, Advances in Neural Information Processing Systems, (2008).
- [47] A.J. ROTHMAN, P.J. BICKEL, E. LEVINA, AND J. ZHU, *Sparse permutation invariant covariance estimation*, Electronic Journal of Statistics, 2 (2008), pp. 494–515.
- [48] K. SCHEINBERG, S. MA, AND D. GOLDFARB, *Sparse inverse covariance selection via alternating linearization methods*, in Proceedings of the Neural Information Processing Systems (NIPS), 2010.
- [49] K. SCHEINBERG AND I. RISH, *Sinco - a greedy coordinate ascent method for sparse inverse covariance selection problem*, (2009). Preprint available at http://www.optimization-online.org/DB_HTML/2009/07/2359.html.
- [50] M. TAO AND X. YUAN, *Recovering low-rank and sparse components of matrices from incomplete and noisy observations*, SIAM J. Optim., 21 (2011), pp. 57–81.
- [51] R. TIBSHIRANI, *Regression shrinkage and selection via the lasso*, Journal of the Royal Statistical Society. Series B, 58 (1996), pp. 267–288.
- [52] C. WANG, D. SUN, AND K.-C. TOH, *Solving log-determinant optimization problems by a Newton-CG primal proximal point algorithm*, preprint, (2009).
- [53] Y. WANG, J. YANG, W. YIN, AND Y. ZHANG, *A new alternating minimization algorithm for total variation image reconstruction*, SIAM Journal on Imaging Sciences, 1 (2008), pp. 248–272.
- [54] Z. WEN, D. GOLDFARB, AND W. YIN, *Alternating direction augmented Lagrangian methods for semidefinite programming*, Mathematical Programming Computation, 2 (2010), pp. 203–230.
- [55] A. WILLE, P. ZIMMERMANN, E. VRANOVÁ, A. FÜRHOLZ, O. LAULE, S. BLEULER, L. HENNIG, A. PRELIC, P. VON ROHR, L. THIELE, ET AL., *Sparse graphical Gaussian modeling of the isoprenoid gene network in Arabidopsis thaliana*, Genome Biology, 5 (2004), p. R92.
- [56] D.M. WITTEN, J.H. FRIEDMAN, AND N. SIMON, *New insights and faster computations for the graphical lasso*, Journal of Computational and Graphical Statistics, 20 (2011), pp. 892–900.
- [57] J. YANG AND Y. ZHANG, *Alternating direction algorithms for l_1 problems in compressive sensing*, SIAM Journal on Scientific Computing, 33 (2011), pp. 250–278.
- [58] M. YUAN, *High Dimensional Inverse Covariance Matrix Estimation via Linear Programming*, The Journal of Machine

- [59] M. YUAN AND Y. LIN, *Model selection and estimation in the Gaussian graphical model*, Biometrika, 94 (2007), pp. 19–35.
 [60] X. YUAN, *Alternating direction methods for sparse covariance selection*, Journal of Scientific Computing, (2009).

Appendix A. Global Convergence Analysis of PGADM.

In this section, we establish the global convergence result of PGADM (Algorithm 1). This convergence proof is not much different with the one given by [57] for compressed sensing problems. We include the proof here just for completeness.

We introduce some notation first. We define $W = \begin{pmatrix} S \\ L \end{pmatrix}$. We define functions $F(\cdot)$ and $G(\cdot)$ as

$$F(R) := \langle R, \hat{\Sigma}_X \rangle - \log \det R,$$

and

$$G(W) := \alpha \|S\|_1 + \beta \text{Tr}(L) + \mathcal{I}(L \succeq 0).$$

Note that both F and G are convex functions. We also define matrix A as $A = [-I_{p \times p}, I_{p \times p}] \in \mathbb{R}^{p \times 2p}$. Now Problem (2.2) can be rewritten as

$$(A.1) \quad \min F(R) + G(W), \quad \text{s.t. } R + AW = 0,$$

and our PGADM (Algorithm 1) can be rewritten as

$$(A.2) \quad \begin{cases} R^{k+1} &:= \operatorname{argmin}_R F(R) + G(W^k) - \langle \Lambda^k, R + AW^k \rangle + \frac{1}{2\mu} \|R + AW^k\|_F^2 \\ W^{k+1} &:= \operatorname{argmin}_W F(R^{k+1}) + G(W) + \frac{1}{2\tau\mu} \|W - (W^k - \tau A^\top (R^{k+1} + AW^k - \mu \Lambda^k))\|_F^2 \\ \Lambda^{k+1} &:= \Lambda^k - (R^{k+1} + AW^{k+1})/\mu. \end{cases}$$

Before we prove the global convergence result, we need to prove the following lemma.

LEMMA A.1. *Assume that (R^*, W^*) is an optimal solution of (A.1) and Λ^* is the corresponding optimal dual variable associated with the equality constraint $R + AW = 0$. Assume the step size τ of the proximal gradient step satisfies $0 < \tau < 1/2$. Then there exists $\eta > 0$ such that the sequence (R^k, W^k, Λ^k) produced by (A.2) satisfies*

$$(A.3) \quad \|U^k - U^*\|_H^2 - \|U^{k+1} - U^*\|_H^2 \geq \eta \|U^k - U^{k+1}\|_H^2,$$

where $U^* = \begin{pmatrix} W^* \\ \Lambda^* \end{pmatrix}$, $U^k = \begin{pmatrix} W^k \\ \Lambda^k \end{pmatrix}$ and $H = \begin{pmatrix} \frac{1}{\mu\tau} I_{p \times p} & 0 \\ 0 & \mu I_{p \times p} \end{pmatrix}$, and the norm $\|\cdot\|_H^2$ is defined as $\|U\|_H^2 = \langle U, HU \rangle$ and the corresponding inner product $\langle \cdot, \cdot \rangle_H$ is defined as $\langle U, V \rangle_H = \langle U, HV \rangle$.

Proof. Since (R^*, W^*, Λ^*) is optimal to (A.1), it follows from the KKT conditions that the followings hold:

$$(A.4) \quad 0 \in \partial F(R^*) - \Lambda^*,$$

$$(A.5) \quad 0 \in \partial G(W^*) - A^\top \Lambda^*,$$

and

$$(A.6) \quad 0 = R^* + AW^*.$$

Note that the first-order optimality conditions for the first subproblem (i.e., the subproblem with respect to R) in (A.2) are given by

$$(A.7) \quad 0 \in \partial F(R^{k+1}) - \Lambda^k + \frac{1}{\mu}(R^{k+1} + AW^k - 0).$$

By using the updating formula for Λ^k , i.e.,

$$(A.8) \quad \Lambda^{k+1} = \Lambda^k - (R^{k+1} + AW^{k+1})/\mu,$$

(A.7) can be reduced to

$$(A.9) \quad 0 \in \partial F(R^{k+1}) - \Lambda^{k+1} - \frac{1}{\mu}(AW^{k+1} - AW^k).$$

Combining (A.4) and (A.9) and using the fact that $\partial F(\cdot)$ is a monotone operator, we get

$$(A.10) \quad \langle R^{k+1} - R^*, \Lambda^{k+1} - \Lambda^* + \frac{1}{\mu}(AW^{k+1} - AW^k) \rangle \geq 0.$$

The first-order optimality conditions for the second subproblem (i.e., the subproblem with respect to W) in (A.2) are given by

$$(A.11) \quad 0 \in \partial G(W^{k+1}) + \frac{1}{\mu\tau}(W^{k+1} - (W^k - \tau A^\top(AW^k + R^{k+1} - \mu\Lambda^k))).$$

Using (A.8), (A.11) can be reduced to

$$(A.12) \quad 0 \in \partial G(W^{k+1}) + \frac{1}{\mu\tau}(W^{k+1} - W^k + \tau A^\top(AW^k - AW^{k+1} - \mu\Lambda^{k+1})).$$

Combining (A.5) and (A.12) and using the fact that $\partial G(\cdot)$ is a monotone operator, we get

$$(A.13) \quad \langle W^{k+1} - W^*, \frac{1}{\mu\tau}(W^k - W^{k+1}) - \frac{1}{\mu}A^\top(AW^k - AW^{k+1}) + A^\top\Lambda^{k+1} - A^\top\Lambda^* \rangle \geq 0.$$

Summing (A.10) and (A.13), and using $R^* = -AW^*$ and $R^{k+1} = \mu(\Lambda^k - \Lambda^{k+1}) - AW^{k+1}$, we obtain,

$$(A.14) \quad \frac{1}{\mu\tau}\langle W^{k+1} - W^*, W^k - W^{k+1} \rangle + \mu\langle \Lambda^{k+1} - \Lambda^*, \Lambda^k - \Lambda^{k+1} \rangle \geq \langle \Lambda^k - \Lambda^{k+1}, AW^k - AW^{k+1} \rangle.$$

Using the notation of U^k , U^* and H , (A.14) can be rewritten as

$$(A.15) \quad \langle U^{k+1} - U^*, U^k - U^{k+1} \rangle_H \geq \langle \Lambda^k - \Lambda^{k+1}, AW^k - AW^{k+1} \rangle,$$

which can be further written as

$$(A.16) \quad \langle U^k - U^*, U^k - U^{k+1} \rangle_H \geq \|U^k - U^{k+1}\|_H + \langle \Lambda^k - \Lambda^{k+1}, AW^k - AW^{k+1} \rangle.$$

Combining (A.16) with the identity

$$\|U^{k+1} - U^*\|_H^2 = \|U^{k+1} - U^k\|_H^2 - 2\langle U^k - U^{k+1}, U^k - U^* \rangle_H + \|U^k - U^*\|_H^2,$$

we get

$$\begin{aligned} (A.17) \quad & \|U^k - U^*\|_H^2 - \|U^{k+1} - U^*\|_H^2 \\ &= 2\langle U^k - U^{k+1}, U^k - U^* \rangle_H - \|U^{k+1} - U^k\|_H^2 \\ &\geq \|U^{k+1} - U^k\|_H^2 + 2\langle \Lambda^k - \Lambda^{k+1}, AW^k - AW^{k+1} \rangle. \end{aligned}$$

Let $\xi := \tau + 1/2$, then we know that $2\tau < \xi < 1$ since $0 < \tau < 1/2$. Let $\rho := \mu\xi$. Then from Cauchy-Schwartz inequality we have

$$\begin{aligned} (A.18) \quad 2\langle \Lambda^k - \Lambda^{k+1}, AW^k - AW^{k+1} \rangle &\geq -\rho\|\Lambda^k - \Lambda^{k+1}\|^2 - \frac{1}{\rho}\|AW^k - AW^{k+1}\|^2 \\ &\geq -\rho\|\Lambda^k - \Lambda^{k+1}\|^2 - \frac{1}{\rho}\lambda_{\max}(A^\top A)\|W^k - W^{k+1}\|^2 \\ &= -\rho\|\Lambda^k - \Lambda^{k+1}\|^2 - \frac{2}{\rho}\|W^k - W^{k+1}\|^2, \end{aligned}$$

where the $\lambda_{\max}(A^\top A)$ denotes the largest eigenvalue of matrix $A^\top A$ and the equality is due to the fact that $\lambda_{\max}(A^\top A) = 2$. Combining (A.17) and (A.18) we get

$$\begin{aligned} (A.19) \quad \|U^k - U^*\|_H^2 - \|U^{k+1} - U^*\|_H^2 &\geq \left(\frac{1}{\mu\tau} - \frac{2}{\rho}\right)\|W^k - W^{k+1}\|^2 + (\mu - \rho)\|\Lambda^k - \Lambda^{k+1}\|^2 \\ &\geq \eta\|U^k - U^{k+1}\|_H^2, \end{aligned}$$

where $\eta := \min\{1 - \frac{2\mu\tau}{\rho}, 1 - \frac{\rho}{\mu}\} = \min\{1 - \frac{2\tau}{\xi}, 1 - \xi\} > 0$. This completes the proof. \square

We are now ready to give the main convergence result of Algorithm (A.2).

THEOREM A.2. *The sequence $\{(R^k, W^k, \Lambda^k)\}$ produced by Algorithm (A.2) from any starting point converges to an optimal solution to Problem (A.1).*

Proof. From Lemma A.1 we can easily get that

- (i) $\|U^k - U^{k+1}\|_H \rightarrow 0$;
- (ii) $\{U^k\}$ lies in a compact region;
- (iii) $\|U^k - U^*\|_H^2$ is monotonically non-increasing and thus converges.

It follows from (i) that $\Lambda^k - \Lambda^{k+1} \rightarrow 0$ and $W^k - W^{k+1} \rightarrow 0$. Then (A.8) implies that $R^k - R^{k+1} \rightarrow 0$ and $R^k + AW^k \rightarrow 0$. From (ii) we obtain that, U^k has a subsequence $\{U^{k_j}\}$ that converges to $\hat{U} = (\hat{W}, \hat{\Lambda})$, i.e., $\Lambda^{k_j} \rightarrow \hat{\Lambda}$ and $W^{k_j} \rightarrow \hat{W}$. From $R^k + AW^k \rightarrow 0$ we also get that $R^{k_j} \rightarrow \hat{R} := -A\hat{W}$. Therefore, $(\hat{R}, \hat{W}, \hat{\Lambda})$ is a limit point of $\{(R^k, W^k, \Lambda^k)\}$.

Note that (A.9) implies that

$$(A.20) \quad 0 \in \partial F(\hat{R}) - \hat{\Lambda}.$$

Note also that (A.12) implies that

$$(A.21) \quad 0 \in \partial G(\hat{W}) - A^\top \hat{\Lambda}.$$

(A.20), (A.21) and $\hat{R} + A\hat{W} = 0$ imply that $(\hat{R}, \hat{W}, \hat{\Lambda})$ satisfies the KKT conditions for (A.1) and thus is an optimal solution to (A.1). Therefore, we showed that any limit point of $\{(R^k, W^k, \Lambda^k)\}$ is an optimal solution to (A.1).

To complete the proof, we need to show that the limit point is unique. Let $\{(\hat{R}_1, \hat{W}_1, \hat{\Lambda}_1)\}$ and $\{(\hat{R}_2, \hat{W}_2, \hat{\Lambda}_2)\}$ be any two limit points of $\{(R^k, W^k, \Lambda^k)\}$. As we have shown, both $\{(\hat{R}_1, \hat{W}_1, \hat{\Lambda}_1)\}$ and $\{(\hat{R}_2, \hat{W}_2, \hat{\Lambda}_2)\}$ are optimal solutions to (A.1). Thus, U^* in (A.19) can be replaced by $\hat{U}_1 := (\hat{R}_1, \hat{W}_1, \hat{\Lambda}_1)$ and $\hat{U}_2 := (\hat{R}_2, \hat{W}_2, \hat{\Lambda}_2)$. This results in

$$\|U^{k+1} - \hat{U}_i\|_H^2 \leq \|U^k - \hat{U}_i\|_H^2, \quad i = 1, 2,$$

and we thus get the existence of the limits

$$\lim_{k \rightarrow \infty} \|U^k - \hat{U}_i\|_H = \eta_i < +\infty, \quad i = 1, 2.$$

Now using the identity

$$\|U^k - \hat{U}_1\|_H^2 - \|U^k - \hat{U}_2\|_H^2 = -2\langle U^k, \hat{U}_1 - \hat{U}_2 \rangle_H + \|\hat{U}_1\|_H^2 - \|\hat{U}_2\|_H^2$$

and passing the limit we get

$$\eta_1^2 - \eta_2^2 = -2\langle \hat{U}_1, \hat{U}_1 - \hat{U}_2 \rangle_H + \|\hat{U}_1\|_H^2 - \|\hat{U}_2\|_H^2 = -\|\hat{U}_1 - \hat{U}_2\|_H^2$$

and

$$\eta_1^2 - \eta_2^2 = -2\langle \hat{U}_2, \hat{U}_1 - \hat{U}_2 \rangle_H + \|\hat{U}_1\|_H^2 - \|\hat{U}_2\|_H^2 = \|\hat{U}_1 - \hat{U}_2\|_H^2.$$

Thus we must have $\|\hat{U}_1 - \hat{U}_2\|_H^2 = 0$ and hence the limit point of $\{(R^k, W^k, \Lambda^k)\}$ is unique. \square

We now immediately have the global convergence result for Algorithm 1 for solving Problem (2.2).

COROLLARY A.3. *The sequence $\{(R^k, S^k, L^k, \Lambda^k)\}$ produced by Algorithm 1 from any starting point converges to an optimal solution to Problem (2.2).*